

WHAT IS CLAIMED IS:

1. A method of generating alternatives to words indicative of recognized speech, recognized based on input speech data, comprising:
 - generating a reference path of recognized words based on the input speech data;
 - receiving an operator selection input indicative of a selected portion of the recognized speech for which alternatives are to be generated;
 - calculating boundary conditions for alternatives based on position bounds of a reference subpath corresponding to the selected portion of the recognized speech; and
 - constructing from a hypothesis store, corresponding to the input speech data, alternate subpaths satisfying the boundary conditions.
2. The method of claim 1 wherein calculating the boundary conditions comprises:
 - determining the position bounds of the reference subpath by identifying, as a beginning boundary condition, a beginning time in the reference path of a first boundary word, the first boundary word preceding the selected portion of the recognized speech.

3. The method of claim 2 wherein calculating the boundary conditions comprises:

determining the temporal bounds of the reference subpath by identifying, as an ending boundary condition, an ending time in the reference path of a second boundary word, the second boundary word following the selected portion of the recognized speech.

4. The method of claim 3 wherein the hypothesis store includes a word hypothesis lattice indicative of entries for a plurality of alternate words corresponding to utterances in the input speech data, each entry including a lexical word, a beginning time in the input speech data corresponding to the lexical word and an ending time in the speech data corresponding to the lexical word, and wherein constructing alternate subpaths comprises:

obtaining from the hypothesis lattice instances of the first boundary word that satisfy the beginning boundary condition.

5. The method of claim 4 wherein constructing alternate subpaths comprises:

concatenating to each of the instances of the first boundary word a string of one or more additional words, wherein the string of one or more additional words satisfies the ending boundary condition.

6. The method of claim 5 wherein constructing alternate subpaths comprises:

concatenating to each of the instances of the first boundary word a string of one or more additional words, wherein the string of one or more additional words ends in an instance of the second boundary word that satisfies the ending boundary condition.

7. The method of claim 5 wherein concatenating comprises:

for each instance of the first boundary word, obtaining from the hypothesis lattice one or more instances of a first intermediate word having a beginning time corresponding to an ending time of the first boundary word and having an ending time.

8. The method of claim 7 wherein concatenating further comprises:

continuously concatenating additional words until an alternate path is formed satisfying the boundary conditions, each of the additional words having a beginning time corresponding to an ending time of a preceding word and an ending time corresponding to a beginning time of a subsequent word.

9. The method of claim 1 wherein constructing alternate subpaths comprises:

beginning construction of up to a predetermined number, X, of hypothesis alternate subpaths;
and

assigning each hypothesis alternate subpath a score that is updated as the hypothesis alternate subpath is constructed.

10. The method of claim 9 wherein assigning a score comprises:

computing an actual score for the alternate subpath through a temporal point to which the hypothesis alternate subpath is constructed; and

computing an estimated score, estimating a score associated with a remainder of the hypothesis alternate subpath to be computed.

11. The method of claim 10 wherein computing an actual score comprises:

computing the actual score as a combination of acoustic scores for words in the hypothesis alternate subpath and language model scores for the words in the hypothesis alternate subpath.

12. The method of claim 11 wherein computing an estimated score comprises:

computing the estimated score as a portion of a reference score assigned to the reference path, the portion of the reference score corresponding to an amount of the hypothesis alternate subpath left to be computed in relation to an overall length of the hypothesis alternate subpath.

13. The method of claim 9 wherein constructing alternate subpaths comprises:

after beginning construction of each of the X hypothesis alternate subpaths, beginning construction of a hypothesis alternate subpath only if its score is one of the top X scores; and
discontinuing construction of all hypothesis alternate subpaths having a score not in the top X scores.

14. The method of claim 13 wherein constructing alternate subpaths comprises:

when construction of X hypothesis alternate subpaths has begun, sorting the hypothesis alternate subpaths by score.

15. The method of claim 14 wherein constructing alternate subpaths comprises:

retaining hypothesis alternate subpaths having the top X scores as the alternate subpaths.

16. The method of claim 15 and further comprising:
presenting the alternate subpaths to the user.

17. The method of claim 1 wherein constructing
alternate subpaths comprises:

constructing alternate subpaths that are
lexically different from the reference
subpath.

18. The method of claim 17 wherein, when fewer than a
predetermined number of lexically different alternate
subpaths can be constructed, extending the boundary
conditions and constructing alternate subpaths
satisfying the extended boundary conditions.

19. The method of claim 1 wherein constructing the
alternate subpaths begins in response to the operator
selection input.

20. A method of generating alternate speech
recognitions for speech data input by an application,
comprising:

receiving the speech data at a speech recognition
engine;

generating a reference speech path and a
hypothesis lattice based on the speech data,
the hypothesis lattice representing
alternate recognition paths corresponding to

the speech data, the hypothesis lattice including a plurality of entries, each entry including an instance of a lexical word, a starting point in the speech data of the instance of the lexical word, an ending point in the speech data of the instance of the lexical word, and a score associated with the instance of the lexical word; storing the hypothesis lattice with the application; receiving a user selection input indicative of a portion of the reference speech path to be corrected; in response to the user selection input, accessing the retrieved hypothesis lattice with the speech recognition engine; and constructing, at the engine, alternate subpaths to replace portions of the reference speech path based on the hypothesis lattice.

21. The method of claim 20 and further comprising: presenting the alternate subpaths to the user.
22. The method of claim 20 wherein constructing alternate subpaths comprises:
 - calculating boundary conditions for alternatives based on temporal bounds of a reference subpath corresponding to the selected portion of the reference speech path; and

constructing from the hypothesis lattice
alternate subpaths satisfying the boundary
conditions.

23. The method of claim 22 wherein calculating the
boundary conditions comprises:

determining the temporal bounds of the reference
subpath by identifying, as a beginning
boundary condition, a beginning time in the
reference speech path of a first boundary
word, the first boundary word preceding the
selected portion of the reference speech
path.

24. The method of claim 23 wherein calculating the
boundary conditions comprises:

determining the temporal bounds of the reference
subpath by identifying, as an ending
boundary condition, an ending time in the
reference speech path of a second boundary
word, the second boundary word following the
selected portion of the reference speech
path.

25. A speech recognition system for receiving a
speech input and generating recognition data
indicative of words recognized in the speech data, and
for generating alternates to words in the recognition
data, comprising:

a decoder generating a reference path including likely words in the speech data and a hypothesis lattice indicative of hypotheses generated based on the speech data; and an alternative generator, coupled to the decoder, configured to receive the reference path and a user selected portion of the reference path to be changed, and to calculate boundary conditions for the selected portion of the reference path to obtain a reference subpath and access the hypothesis lattice to generate alternative subpaths to replace the reference subpath based on the boundary conditions calculated.

26. The speech recognition system of claim 25 wherein the alternative generator comprises:

a boundary calculator configured to receive the user selected portion of the reference path and calculate the boundary conditions for alternative subpaths based on a location of the reference subpath in the reference path.

27. The speech recognition system of claim 26 wherein the alternative generator comprises:

a plurality of construction components coupled to the boundary calculator and configured to construct the alternate subpaths by obtaining hypotheses from the hypothesis

lattice that satisfy the boundary conditions.

28. The speech recognition system of claim 27 wherein the boundary calculator is configured to determining position bounds of the reference subpath in the reference path by identifying, as a beginning boundary condition, a beginning position in the reference path of a first boundary word, the first boundary word preceding the user selected portion of the reference path.

29. The speech recognition system of claim 28 wherein the boundary calculator is configured to determine the position bounds of the reference subpath by identifying, as an ending boundary condition, an ending position in the reference path of a second boundary word, the second boundary word following the user selected portion of the reference path.

30. The speech recognition system of claim 29 wherein the hypothesis lattice includes a word hypothesis lattice indicative of entries for a plurality of alternate words corresponding to utterances in the speech data, each entry including a lexical word, a beginning position in the input speech data corresponding to the lexical word and an ending position in the speech data corresponding to the lexical word.

31. The speech recognition system of claim 30 wherein construction components are configured to obtaining from the hypothesis lattice instances of the first boundary word that satisfy the beginning boundary condition.

32. The speech recognition system of claim 31 wherein the construction components include:

a path extender configured to generate an array of word records forming a concatenation to each of the instances of the first boundary word a string of one or more additional words, wherein the string of one or more additional words satisfies the ending boundary condition.

33. The speech recognition system of claim 32 wherein each of the word records includes an indication of the additional word represented by the word record, preceding word data indicative of preceding words such that a language model score can be calculated, and a pointer to a preceding word record in the array forming part of the alternate subpath being constructed.

34. The speech recognition system of claim 31 wherein the construction components comprise:

a path extender configured to generate an array of word records forming a concatenation to each of the instances of the first boundary word a string of one or more additional words, wherein the string of one or more additional words ends in an instance of the second boundary word that satisfies the ending boundary condition.

35. The speech recognition system of claim 32 wherein the path extender is configured to generate the array of word records by, for each instance of the first boundary word, obtaining from the hypothesis lattice one or more instances of a first intermediate word having a beginning time corresponding to an ending time of the first boundary word and having an ending time.

36. The speech recognition system of claim 35 wherein the path extender is configured to continuously concatenate additional words until an alternate subpath is formed satisfying the boundary conditions, each of the additional words having a beginning time corresponding to an ending time of a preceding word and an ending time corresponding to a beginning time of a subsequent word.

37. The speech recognition system of claim 36 wherein the construction components further include a scoring

component configured to assigning each alternate subpath a score that is updated as the alternate subpath is constructed and wherein the path extender is configured to begin construction, in response to receiving the user selected portion of the reference path, of up to a predetermined number, X, of alternate subpaths.

38. The speech recognition system of claim 37 wherein the scoring component is configured to compute an actual score for the alternate subpath through a position to which the alternate subpath is constructed, and compute an estimated score, estimating a score associated with a remainder of the alternate subpath to be computed.

39. The speech recognition system of claim 38 wherein the path extender is configured to, after beginning construction of each of the X alternate subpaths, continuing construction of an alternate subpath only if its score is one of the top X scores, and discontinuing construction of all alternate subpaths having a score not in the top X scores.

40. The speech recognition system of claim 25 wherein the alternative generator is configured to construct alternate subpaths that are lexically different from the reference subpath, and when fewer than a predetermined number of lexically different alternate

subpaths can be constructed, extending the boundary conditions and constructing alternate subpaths satisfying the extended boundary conditions.

41. The speech recognition system of claim 25 wherein the alternative generator is configured to construct the alternate subpaths in response to the operator selection input.